

Shannon Entropy as an Intrinsic Target Property:  
Toward a Reductionist Model  
of  
Anomalous Cognition

by

Edwin C. May, Ph.D.  
S. James P. Spottiswoode (Consultant)  
and  
Christine L. James  
Science Applications International Corporation  
Cognitive Sciences Laboratory  
Menlo Park, CA

**Abstract**

We propose that the average total change of Shannon's entropy is a candidate for an *intrinsic* target property. We analyze the results of two lengthy experiments that were conducted from 1992 through 1993 and find a significant correlation (Spearman's  $\rho = 0.337$ ,  $df = 31$ ,  $t = 1.99$ ,  $p \leq 0.028$ ) with an absolute measure of the quality of the anomalous cognition. The 1993 result replicated the similar finding from the 1992 study. We describe the methodology, the calculations, and correlations in detail and provide guidelines for those who may wish to conduct similar studies. In addition, we provide circumstantial evidence which leads us toward a reductionist view of anomalous cognition.

## Introduction

The psychophysical properties of the five known senses are well known (Reichert, 1992). At the "front end," they share similar properties. For example, each system possesses receptor cells that convert some form of energy (e.g., photons for the visual system, sound waves for the audio system) into electrochemical signals. The transfer functions are sigmoid; that is, there is a threshold for physical excitation, a linear region, and a saturation level above which more input produces that same output. How these psychophysical reactions translate to sensational experience is not well understood, but all the systems do possess an awareness threshold similar to the subliminal threshold for the visual system.

Since all the known senses appear to share these common properties, it is reasonable to expect that if anomalous cognition (*AC*)<sup>\*</sup> is mediated through some additional "sensory" system, then it, too, should share similar properties. For example, a putative *AC* sensory system should possess receptor cells that have a sigmoidal transfer function and exhibit threshold and saturation phenomena. As far as we know, there are no candidate neurons in the peripheral systems whose functions are currently not understood. So, if receptor cells exist, it is likely that they will be found in the central nervous system. Since 1989, our laboratory has been conducting a search for such receptor sites (May, Luke, Trask, and Frivold, 1990); that activity continues.

There is a second way in which receptor-like behavior might be seen in lieu of a neurophysiology study. If either an energy carrier for *AC* or something that correlated with it were known, then it might be possible to infer sigmoidal functioning at the behavioral level as opposed to the cellular level. Suppose we could identify an intrinsic target property that correlated with *AC* behavior. Then, by manipulating this variable, we might expect to see a threshold at low magnitudes and saturation at high magnitudes.

To construct such an experiment, it is mandatory that we eliminate, as much as possible, all extraneous sources of variance and adopt an absolute measure for the *AC* behavior (Lantz, Luke, and May, 1994). We can reduce one source of variance by considering what constitutes a good target in an *AC* experiment. Delanoy (1988) reported on a survey of the literature for successful *AC* experiments and categorized the target material according to perceptual, psychological and physical characteristics. Except for trends related to dynamic, multi-sensory targets, she was unable to observe systematic correlations of *AC* quality with her target categories.

Watt (1988) examined the target question from a theoretical perspective. She concluded that the "best" *AC* targets are those that are meaningful, have emotional impact, and contain human interest. Those targets that have physical features that stand out from their backgrounds or contain movement, novelty, and incongruity are also good targets.

In trying to understand these findings and develop a methodology for target selection for process-oriented research, we have constructed a metaphor. Figure 1 shows three conceptual domains that contribute to the variability in *AC* experiments. The engineering metaphor of source, transmission, and detector allows us to assign known contributors to the variance of specific domains. Without controlling

\* The Cognitive Sciences Laboratory has adopted the term *anomalous mental phenomena* instead of the more widely known *psi*. Likewise, we use the terms *anomalous cognition* and *anomalous perturbation* for *ESP* and *PK*, respectively. We have done so because we believe that these terms are more naturally descriptive of the observables and are neutral in that they do not imply mechanisms. These new terms will be used throughout this paper.

or understanding these sources, interpreting the results from process-oriented research is problematical, if not impossible.

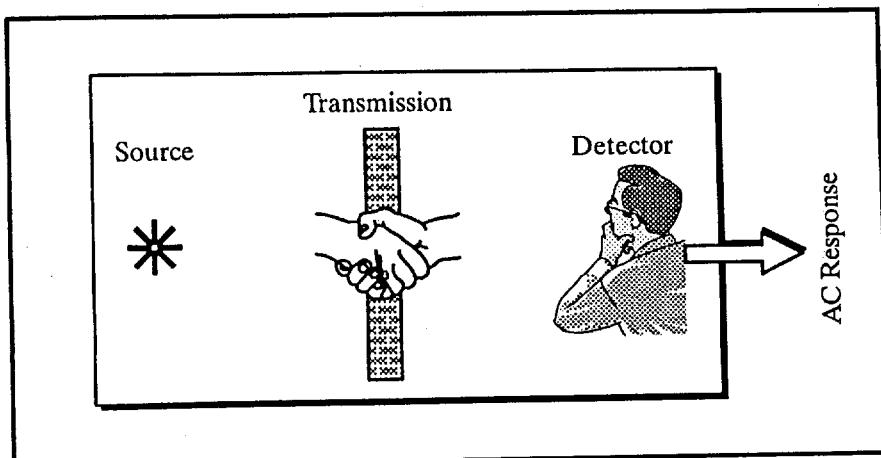


Figure 1. Information-transfer Metaphor

For example, suppose that the quality of an *AC* response actually depended upon the physical size of a target, and that affectivity was also a contributing factor. That is, a large target that was emotionally appealing was reported more often more correctly. Obviously, both factors are important in optimizing the outcome; however, suppose we were studying the effect of target size alone. Then an "attractive" small target might register as well as a less attractive large target and the size dependency would be confounded in unknown ways.

Our metaphor allows us to assign variables, such as these, to specific elements. Clearly, an individual's psychological response to a target is not an *intrinsic* property of a target; rather, it is a property of the receiver. Likewise, size is a physical property of the target and is unrelated to the receiver. Generally, this metaphor allows us to lump together the psychology, personality, and physiology of the receiver and consider these important factors as contributors to a detector "efficiency." If it is true that an emotionally appealing target is easier to sense by some individuals, we can think of them as more efficient at those tasks. In the same way, all physical properties of a target are *intrinsic* to the target and do not depend on the detector efficiency. Perhaps, temporal and spatial distance between target and receiver are intrinsic to neither the target nor the receiver, but rather to the transmission mechanism, whatever that may be.

More than just nomenclature, our metaphor can guide us in designing experiments to decrease certain variabilities in order to conduct meaningful process-oriented research. Some of the methodological improvements seem obvious. If the research objective is to understand the properties of *AC* rather than understanding how an *AC* ability may be distributed in the population, then combining results across receivers should be done with great caution. To understand how to increase high jumping ability, for example, it makes no sense to use a random sample from the general population as high jumpers; rather, find a good high jumper and conduct vertical studies (no pun intended). So, too, is it true in the study of *AC*. We can easily reduce the variance by asking given receivers to participate in a large number of trials and not combining their results.

May, Spottiswoode, and James (1994) suggest that by limiting the number of cognitively differentiable elements within a target, the variance can also be decreased. A further reduction of potential variance can be realized if the target pool is such that a receiver's emotional/psychological response is likely to be more uniform across targets (i.e., reducing the detector variance as shown in Figure 1).

Having selected experienced receivers and attended to these methodological considerations, we could then focus our attention on examining *intrinsic* target properties. If we are successful at identifying one such property, then all process-oriented AC research would be significantly improved because we would be able to control a source of variance that is target specific. The remainder of the paper describes two lengthy studies that provide the experimental evidence to suggest that the average of the total change of Shannon's entropy is one such intrinsic target property.

## Approach

The AC methodological details for the two experiments can be found in Lantz, Luke, and May (1994). In this section we focus on the target calculations and the analysis techniques.

### Target Calculations

Because of the analogy with other sensorial systems, we expected that the change of entropy would be more sensitive than would be the entropy alone. The target variable that we considered, therefore, was the average total change of entropy. In the case of image data, the entropy is defined as:

$$S_k = - \sum_{m=0}^{N_k} p_{mk} \log_2(p_{mk}), \quad (1)$$

where  $p_{mk}$  is the probability of finding image intensity  $m$  of color  $k$ . In a standard, digitized, true color image, each pixel (i.e., picture element) contains eight binary bits of red, green, and blue intensity, respectively. That is,  $N_k$  is 255 (i.e.,  $2^8 - 1$ ) for each  $k$ ,  $k = r, g, b$ . For color,  $k$ , the total change of the entropy in differential form is given by:

$$dS_k = |\nabla S_k| \cdot \vec{dr} + \left| \frac{\partial S_k}{\partial t} \right| dt. \quad (2)$$

We must specify the spatial and temporal resolution before we can compute the total change of entropy for a real image. Henceforth, we drop the color index,  $k$ , and assume that all quantities are computed for each color and then summed.

To compute the entropy from Equation 1, we must specify empirically the intensity probabilities,  $p_m$ . In Lantz, Luke, and May's 1993 experiment, the targets were all video clips that met the following criteria:

- Topic homogeneity. The photographs contained outdoor scenes of settlements (e.g., villages, towns, cities, etc.), water (e.g., coasts, rivers and streams, waterfalls, etc.), and topography (e.g., mountains, hills, deserts, etc.).
- Size homogeneity. Target elements are all roughly the same size. That is, there are no size surprises such as an ant in one photograph and the moon in another.
- Affectivity homogeneity. As much as possible, the targets included materials which invoke neutral affectivity.

For static targets, a single characteristic frame from a video segment was digitized (i.e.,  $640 \times 480$  pixels) for eight bits of information of red, green, and blue intensity. The video image conformed to the NTSC

standard aspect ratio of  $4 \times 3$ , so we arbitrarily assumed an area (i.e., macro-pixel) of  $16 \times 12 = 192$  pixels from which we calculated the  $p_m$ . Since during the feedback phase of a trial the images were displayed on a Sun Microsystems standard 19-inch color monitor, and since they occupied an area approximately  $20 \times 15$  cm square, the physical size of the macro-pixels was approximately 0.5 cm square. Since major cognitive elements were usually not smaller than this, this choice was reasonable—192 pixels were sufficient to provide a smooth estimate of the  $p_m$ .

For this macro-pixel size, the target frame was divided into a  $40 \times 40$  array. The entropy for the  $(i,j)$ 'th macro-pixel was computed as:

$$S_{ij} = - \sum_{m=0}^{N-1} p_m \log_2(p_m),$$

where  $p_m$  is computed empirically only from the pixels in the  $(i, j)$  macro-pixel and  $m$  is the pixel intensity. For example, consider the white square in the upper left portion of the target photograph shown in Figure 2.

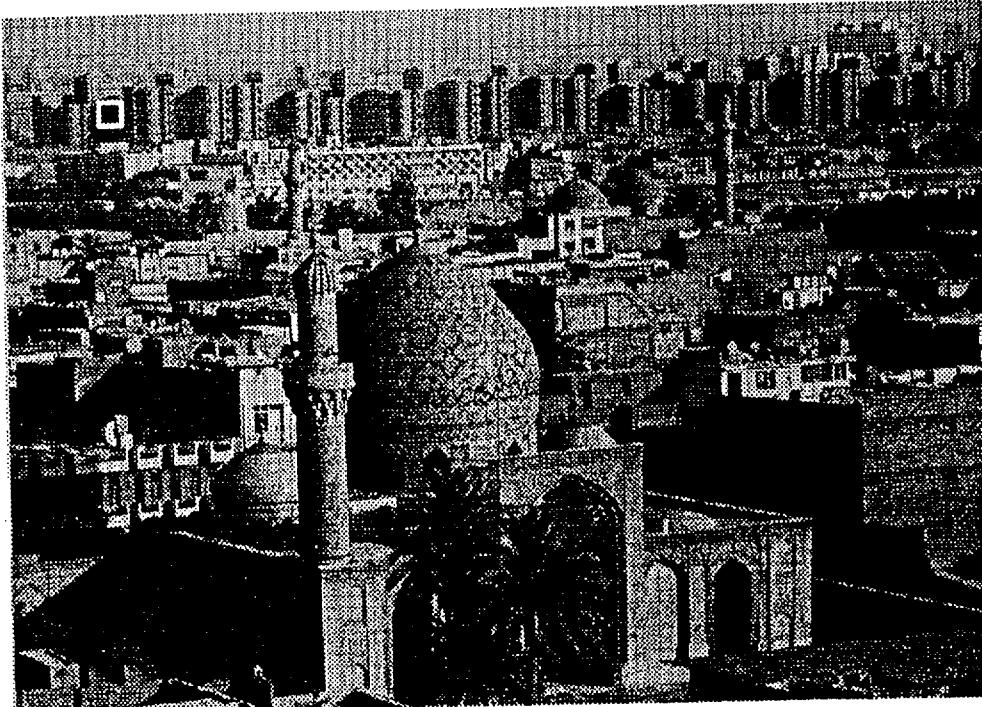


Figure 2. City with a Mosque

The green probability distribution for this macro-pixel (3,3) is shown in Figure 3. The probability density and the photograph itself indicate that most of the intensity in this macro-pixel is near zero (i.e., no intensity of green in this case). In a similar fashion, the  $S_{ij}$  are calculated for the entire scene. Since  $i$  and  $j$  range from zero to 40, each frame contains a total of 1,600 macro-pixels.

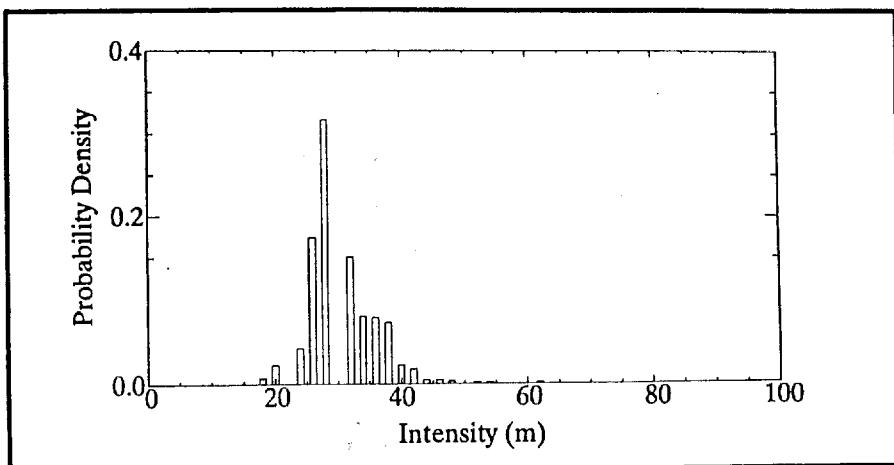


Figure 3. Green Intensity Distribution for the City Target (Macro-pixel 3,3).

We used a standard image processing algorithm to compute the 2-dimensional spatial gradient for each of the 1,600 macro-pixels. The first term in Equation 2 was approximated by its average value over the image.

The total change of entropy for the dynamic targets was calculated in much the same way. The video segment was digitized at one frame per second. The spatial term of Equation 2 was computed exactly as it was for the static frames. The second term, however, was computed from differences between adjacent, 1-second frames for each macro-pixel. Or,

$$\frac{\partial S_{ij}}{\partial t} \approx \frac{\Delta S_{ij}(t)}{\Delta t} = \left| \frac{S_{ij}(t + \Delta t) - S_{ij}(t)}{\Delta t} \right|, \quad (3)$$

where  $\Delta t$  is one over the digitizing frame rate. We can see immediately that the dynamic targets will have a larger  $\Delta S$  than do the static ones because Equation 3 is identically zero for all static targets.

In Lantz, Luke, and May's 1992 experiment, the static targets were digitized from scanned photographs. This difference and its consequence will be discussed below.

### AC-Data Analysis

Rank-order analysis in Lantz, Luke, and May's (1994) experiment demonstrated significant evidence for AC; however, this procedure does not usually indicate the absolute quality of the AC. For example, a response that is a near-perfect description of the target receives a rank of *one*. But a response which is barely matchable to the target may also receive a rank of *one*. Table 1 shows the rating scale that we used to assess the quality of the AC responses, regardless of their rank.

To apply this subjective scale to an AC trial, an analyst begins with a score of *seven* and determines if the description for that score is correct. If not, then the analyst tries a score of *six* and so on. In this way the scale is traversed from *seven* to *zero* until the score-description seems reasonable for the trial.

Table 1.  
0-7 Point Assessment Scale

Score	Description
7	Excellent correspondence, including good analytical detail, with essentially no incorrect information
6	Good correspondence with good analytical information and relatively little incorrect information.
5	Good correspondence with unambiguous unique matchable elements, but some incorrect information.
4	Good correspondence with several matchable elements intermixed with incorrect information.
3	Mixture of correct and incorrect elements, but enough of the former to indicate receiver has made contact with the site.
2	Some correct elements, but not sufficient to suggest results beyond chance expectation.
1	Little correspondence.
0	No correspondence.

## Anomalous Cognition Experiment – 1992

In Lantz, Luke and May's 1992 experiment there were no significant interactions between target condition (i.e., static vs dynamic) and sender condition (i.e., sender vs no sender); therefore, they combined the data for static targets regardless of the sender condition (i.e., 100 trials). The sum-of-ranks was 265 (i.e., exact sum-of-rank probability of  $p \leq 0.007$ , effect size = 0.248). The total sum-of-ranks for the dynamic targets was 300 (i.e.,  $p \leq 0.50$ , effect size = 0.000).

### Entropy Analysis

To examine the relationship of entropy to *AC*, two analysts independently rated all 100 trials (i.e., 20 each from five receivers) from the static-target sessions using the *post hoc* rating scale shown in Table 1. All differences of assignments were verbally resolved, thus the resulting scores represented a reasonable estimate of the visual quality of the *AC* for each trial.

We had specified, in advance, for the correlation with the change of target entropy, we would only use the section of the *post hoc* rating scale that represented definitive, albeit subjective, *AC* contact with the target (i.e., scores four through seven). Figure 4 shows a scatter diagram for the *post hoc* rating and the associated  $\Delta S$  for the 28 trials with static targets that met this requirement. Shown also is a linear least-squares fit to the data and a Spearman rank-order correlation coefficient ( $\rho = 0.452$ ,  $df = 26$ ,  $t = 2.58$ ,  $p \leq 7.0 \times 10^{-3}$ ).

This strong correlation suggests that  $\Delta S$  is an intrinsic property of a static target and that the quality of an *AC* response will be enhanced for targets with large  $\Delta S$ . It is possible, however, that this correlation might be a result of  $\Delta S$  and the *post hoc* rating independently correlating with the targets' visual complexity. For example, an analyst is able to find more matching elements (i.e., a higher *post hoc* rating) in a visually complex target than in a visually simple one. Similarly,  $\Delta S$  may be larger for more complex

targets. If these hypotheses were true, the correlation shown in Figure 4 would not support the hypothesis that  $\Delta S$  is an important intrinsic target property for successful AC.

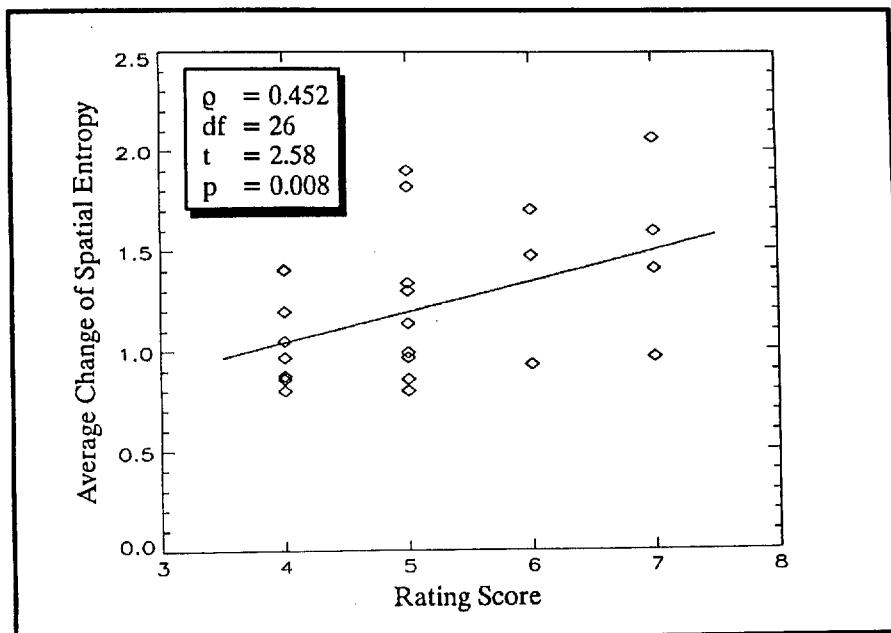


Figure 4. Correlation of *Post Hoc* Score with Static Target  $\Delta S$ .

To check the validity of the correlation, we used a definition of visual complexity, which was derived from a fuzzy set representation of the target pool. We had previously coded by consensus, 131 different potential target elements for their visual impact on each of the targets in the pool. We assumed that the sigma-count (i.e., the sum of the membership values over all 131 visual elements) for each target is proportional to its visual complexity. A description of the fuzzy set technique and a list of the target elements may be found in May, Utts, Humphrey, Luke, Frivold, and Trask (1990).

The Spearman rank correlation between target complexity and *post hoc* rating was small ( $\rho = 0.041$ ,  $df = 98$ ,  $t = 0.407$ ,  $p \leq 0.342$ ). On closer inspection this small correlation was not surprising. While it is true that an analyst will find more matchable elements in a complex target, so also are there many elements that do not match. Since the rating scale (i.e., Table 1) is sensitive to correct and incorrect elements, the analyst is not biased by visual complexity.

Since the change of Shannon entropy is derived from the intensities of the three primary colors (i.e., Equation 1 on page 4) and is unrelated to meaning, which is inherent in the definition of visual complexity, we would not expect a correlation between  $\Delta S$  and visual complexity. We confirmed this expectation when we found a small correlation ( $\rho = -0.028$ ,  $df = 98$ ,  $t = -0.277$ ,  $p \leq 0.609$ ).

Visual complexity, therefore, cannot account for the correlation shown in Figure 4; thus, we are able to suggest that the quality of an AC response depends upon the spatial information (i.e., change of Shannon entropy) in a static target.

A single analyst scored the 100 responses from the dynamic targets using the *post hoc* scale in Table 1. Figure 5 shows the scatter diagram for the *post hoc* scores and the associated  $\Delta S$  for the 24 trials with a

score greater than three for the dynamic targets. We found a Spearman correlation of  $\rho = 0.055$ ,  $df = 22$  ( $t = 0.258$ ,  $p \leq 0.399$ ).

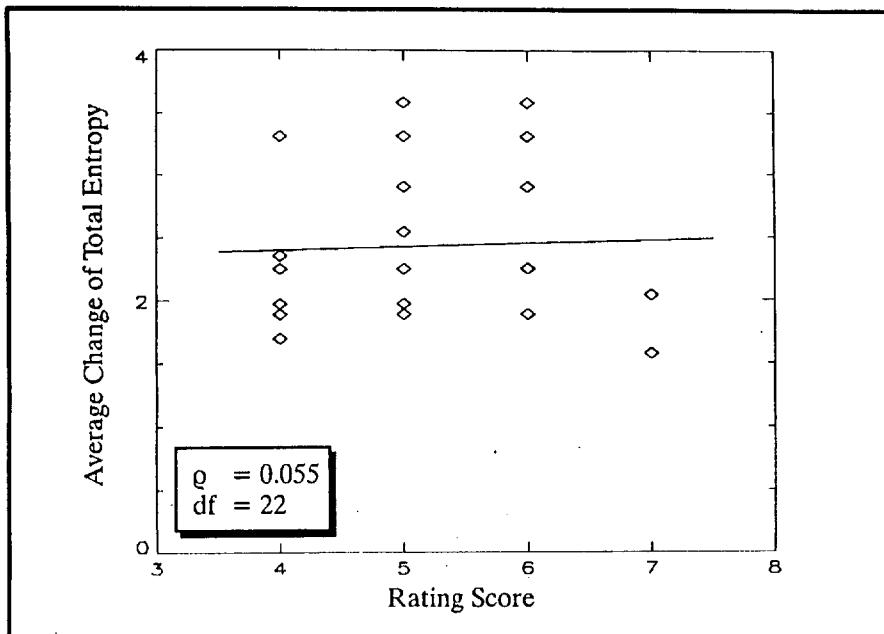


Figure 5. Correlation of *Post Hoc* Score with Dynamic Target  $\Delta S$ .

This small correlation is not consistent with the result derived from the static targets; therefore, we examined this case carefully. The total sum of ranks for the dynamic-target case was exactly mean chance expectation, which indicates that no *AC* was observed (Lantz, Luke, and May, 1994). May, Spottiswoode, and James (1994) propose that the lack of *AC* might be because an imbalance of, what they call, the target pool bandwidth. That is, the number of different cognitive elements in the dynamic pool far exceeded that in the static pool. This imbalance was corrected in the 1993 study and is analyzed below. Regardless, we would not expect to see a correlation if there is no evidence of *AC*.

### Anomalous Cognition Experiment – 1993

The details of the 1993 study may also be found in Lantz, Luke, and May (1994). In that study, they included a static vs dynamic target condition, and all trials were conducted without a sender. They changed the target pools so that their bandwidths were similar. They also included a variety of other methodological improvements, which are not apropos to this discussion.

Lantz, Luke, and May selected a single frame from each dynamic target video clip, which was characteristic of the entire clip, to act as its static equivalent. The static and dynamic targets, therefore, were digitized with the same resolution and could be combined for the correlations.

For each response, a single analyst conducted a blind ranking of five targets—the intended one and four decoys—in the usual way. Lantz, Luke, and May computed effect sizes in the same way as in the 1992 study.

Three receivers individually participated in 10 trials for each target type and a fourth participated in 15 trials per target type. Lantz, Luke, and May reported a total average rank for the static targets of 2.22

for 90 trials for an effect size of 0.566 ( $p \leq 7.5 \times 10^{-5}$ ); the exact same effect size was reported for the dynamic targets.

### Entropy Analysis

Differing from the 1992 experiment, an analyst, who was blind to the correct target choice used the scale, which is shown in Table 1, to assess each response to the same target pack that was used in the rank-order analysis. The average total change of Shannon's entropy (i.e., Equation 2) was calculated for each target as described above. Figure 6 shows the correlation of the blind rating score with this gradient. The squares and diamonds indicate the data for static and dynamic targets, respectively.

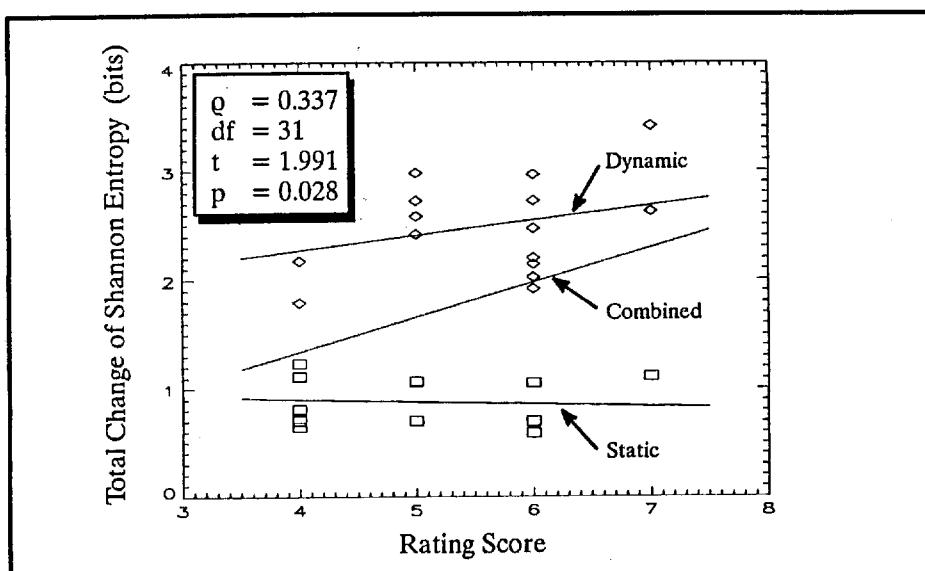


Figure 6. Correlations for Significant Receivers

The key indicates the Spearman correlation for the static and dynamic targets combined. In addition, since the hypothesis was that anomalous cognition would correlate with the total change of the Shannon entropy, Figure 6 only shows the scores in the upper half of the scale in Table 1 for the 70 trials of the three independently significant receivers. The static target correlation was negative ( $\rho = -0.284, df = 13, t = -1.07, p \leq 0.847$ ) and the correlation from the dynamic targets was positive ( $\rho = 0.320, df = 16, t = 1.35, p \leq 0.098$ ). The strong correlation for the combined data arises primarily from the entropic difference between the static and dynamic targets.

### General Conclusions

To understand the differences between the results in the two experiments, we re-digitized the static set of targets from the 1992 experiment with the same hardware and software that was used in the 1993 study. With this new entropy data, the correlation dropped from a significant 0.452 to 0.298 which is not significant ( $t = 1.58, df = 26, p \leq 0.063$ ). Combining this data with the static results from the 1993 experiment (i.e., significant receivers) the static correlation was  $\rho = 0.161, df = 41 (t = 1.04, p \leq 0.152)$ . The correlation for the static targets from the 1992 experiment combined with the significant static and dynamic data from the 1993 experiment was significant ( $\rho = 0.320, df = 59, t = 2.60, p \leq 0.006$ ). These *post hoc* results are shown in Figure 7. The combined data from the two experiments, including all re-

ceivers and all scores greater than four, give a significant correlation ( $\rho = 0.258, df = 64, t = 2.13, p \leq 0.018$ ).

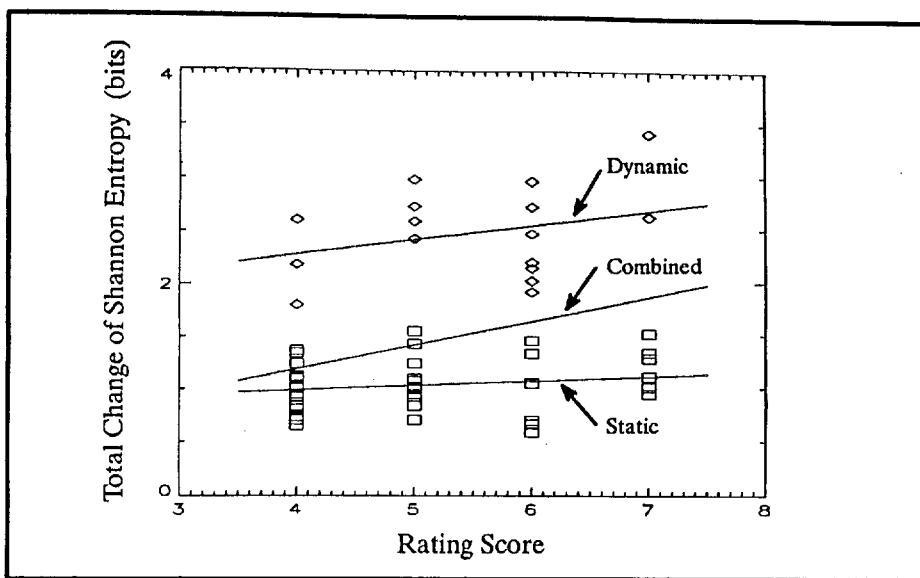


Figure 7. Correlations for Combined Experiments

We conclude that the quality of *AC* appears to correlate linearly with the average total change of the Shannon entropy, which is an *intrinsic* target property.

These two experiments may raise more questions than they answer. If our conservative approach, which assumes that *AC* functions similarly to the other sensorial systems, is correct, we would predict that the *AC* correlation with the frame entropy should be smaller than that for the average total change of the entropy. We computed the total frame entropy from the  $p_i$  all of the  $640 \times 480$  pixels. The resulting correlation for the significant receivers in the 1993 experiment was  $\rho = 0.234, df = 31 (t = 1.34, p \leq 0.095)$ . This correlation is considerably smaller than that from the gradient approach, however, not significantly so. We computed the average of the  $S_{ij}$  for the 1,600 macro-pixels as a second way of measuring the spatial entropic variations. We found a significant Spearman's correlation of  $\rho = 0.423, df = 31 (t = 2.60, p \leq 0.007)$  for the significant receivers in the 1993 experiment. The difference between the correlation of the quality of the *AC* with the frame entropy and with either measure of the spatial gradient is not significant; however, these large differences are suggestive of the behavior of other sensorial systems (i.e., an increased sensitivity with change of the input).

We have quoted a number of different correlations under varying circumstances and have labeled these as *post hoc*. For example, hardware limitations in 1992 prevented us from combining those data with the data from 1993. Thus, we recalculated the entropies with the upgraded hardware in 1993 and recomputed the correlations. Our primary conclusions, however, are drawn only from the static results from the 1992 experiment and the confirmation from the combined static and dynamic 1993 results.

It is clear from our analysis that we may have identified an intrinsic target property that correlates with the quality of anomalous cognition. Our results suggest a host of new experiments and analyses before we can come to this conclusion with certainty. For example, suppose we construct a new target pool that is maximized for the gradient of Shannon's entropy yet meets reasonable criteria for the target pool

bandwidth. If the Shannon information is important, than we should see exceptionally strong *AC*. We also must improve the absolute measure of *AC*. While dividing our zero-to-seven rating scale in two makes qualitative sense, it was an arbitrary decision. Rank order statistics are not as sensitive to correlations as are absolute measures (Lantz, Luke, and May, 1994); but, perhaps, if the *AC* effect size is significantly increased with a proper target pool, the rank-order correlations will be strong enough. It may be time consuming; however, it is also important to understand the dependency of the correlation on the digitizing resolution. In the first experiment, we digitized the hard copy photographs using a flatbed scanner with an internal resolution of 100 dots/inch and used  $640 \times 480$  pixels for the static and dynamic targets in the second experiment. Why did the correlation drop for the static targets by nearly 35 percent when the digitizing resolution decreased by 20 percent?

We noticed, *post hoc*, that the correlations exhibit large oscillations around zero below the cutoff score of four. If we assume there is a linear relationship between *AC* scores and the total change of Shannon entropy, we would expect unpredictable behavior for the correlation at low scores because they imply chance matches with the target and do not correlate with the entropy.

Since we are suggesting a reductionist perspective, we speculate that the linear correlation suggests behavioral, albeit circumstantial, evidence for receptor-like functioning for the detection of *AC*. To determine if this is true, we must identify threshold and saturation limits.

It is absolutely critical to confirm our overall results and to provide answers to some of the enigmas from our experiment. If we have identified an *intrinsic* target property, then all of our research can precede more efficiently. Consider the possibilities if we were able to construct a target pool and eliminate a known source of variance. Psychological and physiological factors would be much easier to detect. Given the availability of inexpensive video digitizing boards for personal computers, replication attempts are easily within the grasp of research groups with modest operating budgets.

## References

Bem, D. J. and Honorton, C. (1994). Does psi exist? Replicable evidence for an anomalous process of information transfer. *Psychological Bulletin*, 115, No. 1, 4-18.

Delanoy, D. L. (1988), Characteristics of successful free-response targets: Experimental findings and observations. *Proceedings of Presented Papers*, The Parapsychological Association 31st Annual Convention, Montreal, Canada, 230-246.

Watt, C. (1988). Characteristics of successful free-response targets: Theoretical considerations. *Proceedings of Presented Papers*, The Parapsychological Association 31st Annual Convention, Montreal, Canada, 247-263.

Lantz, N. D. and Luke, W. L. W., and May, E. C. (1994). Target and sender dependencies in anomalous cognition experiments. Submitted for publication in the *Journal of Parapsychology*.

May, E. C., Luke, W. L. W., Trask, V. V., and Frivold, T. J. (1990). Observation of neuromagnetic fields in response to remote stimuli. *Proceedings of Presented Papers*, The Parapsychological Association 33rd Annual Convention, National 4-H Center, Chevy Chase, MD, 168-185.

May, E. C., Spottiswoode, S. J. P., and James, C. L. (1994). Managing the target pool bandwidth: Noise reduction in anomalous cognition experiments. Submitted for publication in the *Journal of Parapsychology*.

May, E. C., Utts, J. M., Humphrey, B. S., Luke, W. L. W., Frivold, T. J., and Trask, V. V. (1990). Advances in remote-viewing analysis. *Journal of Parapsychology*, 54, 193-228.

Reichert, H. (1992). *Introduction to Neurobiology*, Oxford University Press, New York, NY.